

Magic-state distillation with the four-qubit code

Adam M. Meier,^{1,2,*} Bryan Eastin,³ and Emanuel Knill^{1,2}

¹*University of Colorado, Boulder, CO*

²*National Institute of Standards and Technology, Boulder, CO*

³*Northrop Grumman Corporation, Baltimore, MD*

(Dated: March 2, 2013)

The distillation of magic states is an often-cited technique for enabling universal quantum computing once the error probability for a special subset of gates has been made negligible by other means. We present a routine for magic-state distillation that reduces the required overhead for a range of parameters of practical interest. Each iteration of the routine uses a four-qubit error-detecting code to distill the $+1$ eigenstate of the Hadamard gate at a cost of ten input states per two improved output states. Use of this routine in combination with the 15-to-1 distillation routine described by Bravyi and Kitaev allows for further improvements in overhead.

Many techniques for robustly implementing quantum gates most naturally generate only a finite subset of the unitary operators. Frequently, the naturally convenient quantum operations generate the full set of Clifford operations, which consists of the Clifford group of unitaries augmented by measurement and state preparation in the standard basis. Clifford operations are sufficient for stabilizer-state preparations and measurements and thus underlie stabilizer-based error correction and much of the associated theory of fault tolerance. Though inadequate for universal quantum computing, the Clifford operations can be supplemented by any unitary outside of the Clifford group to obtain a universal set [1]. Consequently, the problem of achieving universality is often reduced to that of finding a way of robustly implementing a single non-Clifford unitary gate.

Given the ability to perform Clifford operations, non-Clifford gates can be indirectly implemented using certain non-stabilizer states as a consumable resource. The advantage of this approach lies in the possibility of distilling such resource states prior to use. Distillation is a technique whereby a collection of independently prepared faulty resource states can be converted into a smaller number of resource states whose fidelity with respect to the ideal state is higher. Some states have the property that one can distill them using only Clifford operations. States that are both sufficient for universality and distillable in this way are known as *magic states*. Magic-state distillation allows faulty magic states to be used as a resource for robust universal quantum computing.

The notion of magic states was introduced by Bravyi and Kitaev [2], who showed that the (magic) eigenstates of the one-qubit Clifford gates T and H can be distilled from copies of these states with error probabilities of up to 0.173 and 0.141 per state, respectively. Their distillation routines work by projecting several such faulty copies of a specified magic state (henceforth, resource state) into a stabilizer code and then decoding the result,

checking for and discarding on any indication of error. Distillation of the T -eigenstate $|T\rangle$ employs a projection onto the 5-qubit distance-3 code, while distillation of the H -eigenstate $|H\rangle$ relies on the 15-qubit Reed-Muller code; both distillation routines result in one improved resource state. We refer to the $|H\rangle$ -distillation routine as the 15-to-1 routine.

An apparently distinct routine for distilling $|H\rangle$ using the 7-qubit Steane code was proposed previously by Knill [3], but Reichardt found the two routines to be equivalent [4]. Reichardt additionally showed that the error threshold for distilling $|H\rangle$ could be improved from 0.141 to 0.146 via a 7-to-1 distillation routine, thereby proving that every faulty $|H\rangle$ outside of the set of stabilizer states is distillable with a finite routine. Campbell and Browne proved the impossibility of a similar result for $|T\rangle$ by showing that no finite distillation routine is capable of distilling faulty $|T\rangle$ arbitrarily near the boundary of the stabilizer states [5, 6].

The focus of each of the aforementioned papers is on the threshold for magic-state distillation, but the efficiency of a distillation routine is crucial to its practical utility. Of particular concern is the number of faulty resource states required as input to distill each resource state of some desired quality. This ratio contributes strongly to the overhead required to implement a quantum computation using magic-state distillation [7], potentially increasing the number of qubits and gates required by a large multiplicative factor. With this in mind, we describe a routine for distilling $|H\rangle$ that reduces the number of input resource states required per output state, distilling 2 improved resource states from 10. The routine can be used either solely or in combination with previously developed routines to obtain resource reductions for a variety of parameter ranges of interest.

After explaining the needed background in Sec. I, we introduce and analyze the proposed distillation routine in Secs. II and III and compare it to the 15-to-1 routine in Sec. V. Sec. IV explains how sequential distillation rounds can be combined. Concluding remarks appear in Sec. VI.

* Adam.Meier@colorado.edu

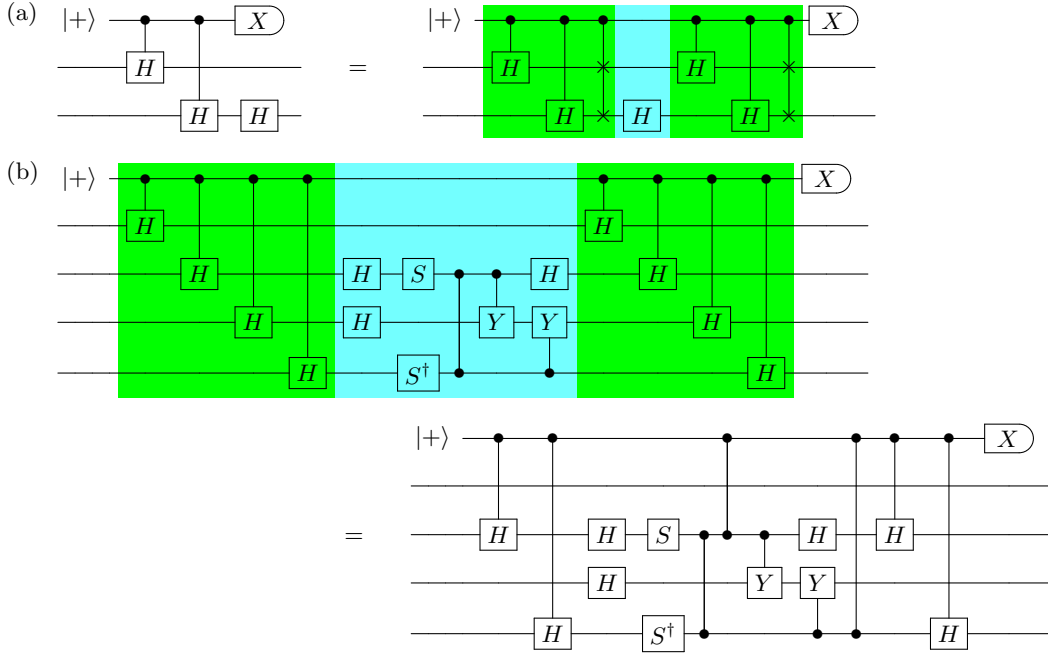


FIG. 2. Circuits that detect whether two input states are in the subspace spanned by $|H\rangle|-H\rangle$ and $|-H\rangle|H\rangle$, when the input states are (a) unencoded and (b) encoded in the four-qubit code. Corresponding blocks of the circuits in parts (a) and (b) are indicated by shading. The translation from (a) to (b) relies on the fact that, for the four-qubit code, transversal Hadamard effects $\bar{H}_1\bar{H}_2\bar{X}_{12}$. The equivalence in b) eliminates four C_H gates, reducing the number of resource states required by eight. In total, this figure shows that only four C_H gates are required to project the two qubits encoded in the four-qubit code into either the subspace spanned by $\{|H\rangle|-H\rangle, |-H\rangle|H\rangle\}$ or that spanned by $\{|H\rangle|H\rangle, |-H\rangle|-H\rangle\}$. The circuits shown also apply an incidental Hadamard gate to the second logical qubit. Further details are given in Fig. 8 in the appendix.

Because $Y|H\rangle = |-H\rangle$, we can characterize any faulty $|H\rangle$ state that has been twirled with \mathcal{H} as suffering from stochastic Y errors with some probability p that depends on the input state ρ . We assume throughout this paper that resource states are twirled prior to use.

We label distillation routines by their input/output ratios, so an m -to- n distillation routine takes m resource states as input and produces n resource states as output.

II. 10-TO-2 DISTILLATION ROUTINE

The basic form of our routine for magic-state distillation is as follows: Resource states are encoded into a quantum code; these encoded resource states are verified through an encoded measurement; and finally the code is decoded, leaving, when no errors are indicated, resource states of better quality. The intuition behind this approach is that one would like simply to measure whether a resource state is good, but doing so requires additional resource states whose own errors might go undetected in such a measurement. Errors on these states are rendered detectable by performing an encoded version of the measurement in a fault-tolerant fashion. This is the approach employed in reference [3] for distilling $|H\rangle$ using the 7-qubit Steane code. The routine described here is instead based on the 4-qubit error-detecting code.

As the $+1$ eigenstate of the Hadamard operator, the

state $|H\rangle$ can be verified by measuring H . Measurement of the Hadamard operator is impossible using only Clifford operations, but it can be accomplished, as shown in Fig. 1, with the help of two additional $|H\rangle$ states.

To render errors during the Hadamard measurement detectable, the routine first encodes a pair of faulty resource states into the \mathcal{C}_4 code [3] and then performs an encoded measurement $\bar{H}_1\bar{H}_2$ on the pair. This measurement determines whether the pair is in the logical subspace spanned by $|H\rangle|-H\rangle$ and $|-H\rangle|H\rangle$ and can therefore detect whether one of the states had an error (see Sec. III).

The \mathcal{C}_4 code is a $[[4, 2, 2]]$ quantum code defined by the stabilizer generator matrix:

$$\begin{bmatrix} X \otimes X \otimes X \otimes X \\ Z \otimes Z \otimes Z \otimes Z \end{bmatrix}. \quad (2)$$

Our choices for logical X and Z operators are:

$$\begin{aligned} \bar{X}_1 &= X \otimes X \otimes I \otimes I, \\ \bar{Z}_1 &= Z \otimes I \otimes I \otimes Z, \\ \bar{X}_2 &= X \otimes I \otimes I \otimes X, \quad \text{and} \\ \bar{Z}_2 &= Z \otimes Z \otimes I \otimes I. \end{aligned}$$

Because any one-qubit Pauli operator anticommutes with some stabilizer generator of \mathcal{C}_4 , it is possible to detect any error on a single qubit of the code.




FIG. 3. Circuit identities for propagating Y errors on gate states. (a) The rule for propagating a Y error on a resource state used to apply a $Y(\pm\pi/4)$ gate as in Fig. 1(c). The effect of the error is the same as applying the Y error after the gate. (b) A Y error on the first resource state required to implement a CH gate using the circuits in Fig. 1 propagates to a Z error on the control and a Y error on the target. An error on the second resource state propagates trivially to a Y error on the target.

The set of stabilizer generators of \mathcal{C}_4 is symmetric with respect to exchange of X and Z , so $H \otimes H \otimes H \otimes H$ is a valid encoded gate, and for the choice of logical Pauli operators given above it effects a logical Hadamard on both encoded qubits followed (or, equivalently, preceded) by a logical SWAP . Consequently, the controlled- $(\bar{H}_1 \bar{H}_2 \bar{X}_{12})$ gate (the control is unencoded and the target is encoded in \mathcal{C}_4) can be accomplished using a sequence of four CH gates. Using this gate one can derive a circuit that implements the encoded measurement, $\bar{H}_1 \bar{H}_2$, as shown in Fig. 2, by means of four CH gates implemented with a total of eight resource states.

The final step of the distillation routine is to use Clifford operations to decode the logical qubits and measure the syndrome of the \mathcal{C}_4 code, leaving two output resource states. The routine succeeds and accepts the output if neither the encoded measurement nor the syndrome indicates an error. Otherwise the output is discarded. We analyze the error patterns for the full distillation circuit, shown in Fig. 4(a), in the next section.

III. ANALYSIS

Given perfect Clifford operations and twirled resource states, the only possible errors in our distillation circuit are Y errors on the input resource states. For simplicity we assume that the input states to be distilled are independent and all have the same error probability p .

As described in the previous section, the ten input resource states can be partitioned into two resource states that are encoded into the code \mathcal{C}_4 (data states) and four pairs of resource states used to implement CH gates (gate states). The effect of one error on either type of resource state can be understood as follows.

A Y error on one of the data states becomes an encoded Y error, which flips the outcome of the encoded measurement (the measurement of $\bar{H}_1 \bar{H}_2$) and is thus detected by

the routine. The decoding exactly reverses the encoding, and the logical gates in between preserve logical Y errors, so errors on data states persist on the output and are not detected by the syndrome measurement.

As shown in Fig. 3, a Y error on one of the gate states causes the intended CH gate to act as CH followed by either $Z \otimes Y$ or $I \otimes Y$, depending on which resource state was in error. Using circuit identities, these errors can be propagated to a common location just before the second set of CH gates, as depicted in Fig. 4(b). At this location, such an error appears as a combination of some logical operator and a Y error on a single qubit, which is not an encoded Pauli operator for the \mathcal{C}_4 code. This Y error is followed only by logical operators, which cannot take an error subspace to a non-error subspace, and decoding, which returns a syndrome indicating whether the state is in an error subspace. Consequently, a single error on a gate state is detected by the syndrome measurement.

The effect of multiple errors is best understood by propagating the errors from both gate and data states to two locations, as described in Fig. 4. The Y Pauli operators from any pair of errors on gate states (described above) combine to form a logical operator for the code, so any even number of errors on such states will fail to be detected by the decoder, while any odd number of errors will be detected. For each error pattern that is not detected by the syndrome, one can consider the effect of the logical errors on the encoded information and encoded measurement. For example, even numbers of Z errors on the encoded-measurement ancilla will cancel and cause the distillation to be accepted. Each pattern of errors on the resource states can then be classified first by whether it is detected and then by whether it causes a non-trivial logical error. Because errors on each state are considered to be equiprobable and independent, this enumeration determines the probability $a(p)$ of the distillation routine accepting and the marginal error probability $e(p)$ of an output state conditional on acceptance. It happens that $e(p)$ does not depend on which of the two output states is considered.

Based on the observation that any single error results in rejection, a simple estimate of the acceptance probability is $a(p) = 1 - 10p + O(p^2)$. The exact accounting yields

$$a(p) = 1 - 10p + 58p^2 - 192p^3 + 400p^4 - 544p^5 + 480p^6 - 256p^7 + 64p^8.$$

The probability of an undetected error on the output states is the probability that the routine accepts and that the output nevertheless has an error. Because any single error is detected, this probability has order p^2 . The marginal undetected-error probability of the first (or identically the second) output state is

$$u(p) = 9p^2 - 56p^3 + 160p^4 - 256p^5 + 240p^6 - 128p^7 + 32p^8.$$

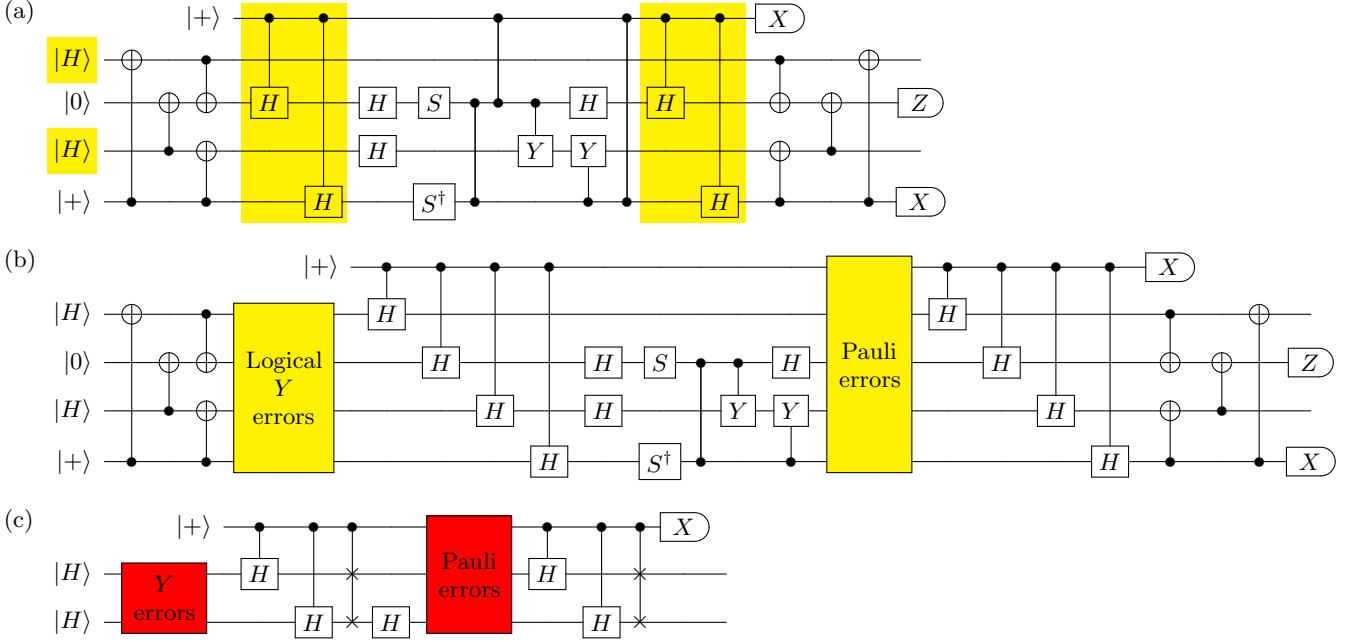


FIG. 4. Illustration of a method of classifying errors in our distillation routine. The first circuit shows the full distillation routine with possible error locations shaded. Using circuit identities, errors at any of these locations can be concentrated into one of two regions (and types) yielding an equivalent circuit of the form shown in (b). Any Pauli errors on the lower four (code) qubits in the second error region of (b) will be detected by the decoding circuit unless they act as a logical (Pauli) operator on the code. The remaining possible errors, those undetectable by the syndrome measurement, may then be enumerated and classified efficiently using the logical circuit shown in (c).

For our purposes, this is the quantity of interest, but one can also compute the probability of at least one undetected error on the two outputs. This is given by

$$u_2(p) = 13p^2 - 80p^3 + 228p^4 - 368p^5 + 352p^6 - 192p^7 + 48p^8.$$

The quality of the distillation routine's output is quantified by the marginal probability of error of an output state conditional on acceptance:

$$e(p) = u(p)/a(p). \quad (3)$$

The corresponding probability of at least one error on the two outputs conditioned on acceptance is $e_2(p) = u_2(p)/a(p)$. It can be shown numerically that $e_2(p) \leq 2e(p) - e(p)^2$, so errors on the two output states are positively correlated. In fact, the probability of an error on both output states is of order p^2 .

IV. DISTILLATION SEQUENCES

The ultimate goal of magic-state distillation is to produce resource states of sufficiently high quality that they can be used to implement all non-Clifford gates in a computation without significantly increasing the probability that the computation will fail. A generic computation will fail if any single gate fails, so the probability of one or more errors on the R resource

states employed in a computation must be much less than 1 to ensure that the computation succeeds with high probability. By the union bound, it is sufficient that the marginal probability of error on each resource state be much less than $1/R$. Strong correlations can reduce this requirement on marginal probabilities, but for independent errors the bound is necessary. Consequently, the proximate goal of magic-state distillation is to produce resource states such that the marginal probability of error for any single state is bounded from above by some goal error probability, $e_g \ll 1/R$. In algorithms currently envisioned for quantum computers, R can easily be 10^{10} or more.

In order to obtain resource states with very low probabilities of error, it is necessary to use multiple rounds of distillation, where the input to each round is produced by the preceding one. We consider a sequence of such rounds where each is based on a single but possibly round-dependent distillation routine. In a round based on an m -to- n distillation routine, the output resource states from the preceding round are grouped into blocks of size m , and each block is then distilled to n states, which may, in general, have correlated errors.

The sequence of rounds is chosen to minimize the number of input resource states needed to produce a given number of output states with marginal probability of error e_g or less. In practice, we are interested in the case where the number of resource states to be prepared is very large, allowing us to consider only the asymptotic

cost. The cost is defined as the number of input resource states used per output resource state produced. For one round of the 10-to-2 distillation routine, the cost is $\frac{10}{2a(p)}$, with marginal probability of error $e(p)$ on the output states conditioned on acceptance.

The one-round expression for marginal probability of error given in Sec. III assumes that the input resource states suffer from errors independently and with equal probability. Generally, however, the output of a distillation routine need not satisfy either restriction, which poses a concern for distillation sequences involving multiple rounds. If necessary, distillation routines can be output symmetrized by randomly permuting the output states, thereby ensuring that the output states from a given round all have the same error probability. Independence is a concern whenever a routine that outputs more than one state per instance is used, since errors on the states output by one instance of such a routine are usually not independent. For example, the probability of two errors in the output of the 10-to-2 distillation routine is of the same order as that for one error. Performing a distillation routine using such correlated states as input can substantially increase the output error probability. To avoid this effect, it is sufficient to ensure that no instance of a routine depends on more than one output from any previously executed instance of a routine. The following lemma and its corollary show that this strategy works without an increase in asymptotic cost. As a consequence we can calculate the asymptotic cost as if the output states of all routines were completely independent.

Lemma IV.1. *Let \mathcal{D} be an m -to- n output-symmetrized distillation routine with acceptance probability $a(p)$ and conditional output error probability $e(p)$ for each output state. Given a block of K independent resource states, each with error probability p , one can produce n blocks of output states where each block's states are independent within the block and have probability of error $e(p)$. Each block contains $a(p)\lfloor K/m \rfloor$ states on average.*

Proof. Partition the K resource states into $\lfloor K/m \rfloor$ sets of m states, discarding any remaining ones. Apply \mathcal{D} to each set of m states, getting n states with probability $a(p)$ in each case. Conditional on acceptance, each output state has marginal probability of error $e(p)$, though these errors are not independent. Form n blocks by taking the j^{th} output state from each successful distillation, for $j = 1, \dots, n$. These blocks have the desired properties for a given pattern of distillation successes. Because the error probabilities of the j^{th} output states of the successful distillations do not depend on the pattern of successes, any such arrangement of the output states that depends only on the pattern of successes preserves this independence. \square

Corollary IV.2. *If, in Lem. IV.1, K is random with average $\langle K \rangle$, then the expected total number of output*




FIG. 5. Plots of the marginal error probabilities conditional on acceptance for the 10-to-2 (solid) and 15-to-1 (dotted) routines. The dashed line indicates the output if no distillation is performed. The thresholds for the two routines are determined by the first intersections with this line.

states is at least $a(p)n \left(\frac{\langle K \rangle}{m} - 1 \right)$. The average size of each of the n output blocks of independent states is at least $a(p) \left(\frac{\langle K \rangle}{m} - 1 \right)$.

A multi-round distillation routine can now be formulated as follows: Assume that after round $l-1$ there are N_{l-1} blocks of resource states, where within each block the states are independent with identical error probabilities p_{l-1} , and the number of states in each block is $\langle K_{l-1} \rangle$ on average. Applying the procedure of Lem. IV.1 to each block with a $m_l \rightarrow n_l$ distillation routine \mathcal{D}_l yields $N_l = N_{l-1}n_l$ output blocks, where each output block has $\langle K_l \rangle \geq a_l(p_{l-1}) \left(\frac{\langle K_{l-1} \rangle}{m_l} - 1 \right)$ resource states on average, independent within a block and each with error probability $p_l = e_l(p_{l-1})$. The first round starts with K_0 independent resource states, each of which suffer an error with probability p_0 . For large K_0 , the constant offsets of -1 in the expressions are negligible. Consequently, the asymptotic cost c_l of resource-state production after round l satisfies $c_l = \frac{m_l}{n_l a_l(p_{l-1})} c_{l-1}$, where $c_0 = 1$. The error probability after round l satisfies $p_l = e_l(p_{l-1})$.

V. COMPARATIVE PERFORMANCE

At present, practical fault-tolerant architectures require physical-gate error probabilities well below .01, which suggests that it should be possible to directly prepare resource states with error probabilities of a few percent or less. Given such states, the most practical routine for $|H\rangle$ distillation developed to date is the 15-to-1 routine. In this section, we compare the performance of the 10-to-2 routine to that of the 15-to-1 routine and consider the effect of using them in concert.




FIG. 6. Output error probability as a function of input error probability for various sequences of the 10-to-2 (A) and 15-to-1 (B) $|H\rangle$ distillation routines. Each curve is labeled by the associated sequence of routines, e.g., *BAA* denotes the 15-to-1 routine followed by two rounds of the 10-to-2 routine. The region directly underneath each labeled curve is the region in which the labeled strategy is preferred.

Routines for magic-state distillation are typically judged on the basis of their threshold, that is, the error probability below which resource states can be successfully distilled. At the threshold, a distillation routine outputs resource states no better than the inputs. Thus, the threshold p_t for the 10-to-2 routine can be determined from Eq. (3) by considering solutions to $p_t = e(p_t)$. This yields a threshold of $p_t = 0.089$, which is substantially below the threshold of 0.141 for the 15-to-1 routine [2], but either threshold should be adequate for the error regime of interest. The curves for the marginal output error probability of the 10-to-2 and 15-to-1 routines are plotted in Fig. 5.

The efficiency of a distillation routine can be characterized, as detailed in Ref. [2], by the output error probability as a function of the number of resource states employed. In the limit of small initial error probability p , the output error probability for the 10-to-2 routine after l rounds of distillation is $\frac{1}{9}(9p)^{2^l}$. In the limit of both small p and many output states, l rounds of distillation require $k = 5^l$ input resource states per output. Consequently, taking l to be continuous, the asymptotic output error probability as a function of the number of input resource states expended is $\frac{1}{9}(9p)^{k^\xi}$, where $\xi = \frac{1}{\log_2(5)} \approx .43$. The corresponding exponent for the 15-to-1 routine is .4, so the 10-to-2 routine performs slightly better for this metric. However, these smooth functions hide the step

discontinuities induced by using sequences of increasing integral lengths (as seen in Fig. 7) and can be misleading for practical comparisons.

Of greater utility to us is the cost, in resource states consumed per output state, required to obtain resource states of sufficiently high quality for useful quantum computations, given resource states with error probabilities in the range of 0.01 to 10^{-5} . The cost depends on the distillation sequence, which generally entails multiple rounds of distillation. For the purpose of optimizing the distillation sequence, one can consider arbitrary routines at each round. Here, we consider sequences involving the 10-to-2 and 15-to-1 distillation routines.

In Fig. 6 we plot the output error probability as a function of input error probability for various sequences. Data for the 15-to-1 routine was computed using the expressions corresponding to $a(p)$ and $e(p)$ in Eq. (35) and Eq. (36) of Ref. [2]. In the region plotted, distillation sequences with higher output error (lower curves) also require fewer input resource states per output state. Consequently, for a given output error goal, e_g , and input error probability p , the label of the nearest curve above the point (p, e_g) in the plot gives the best distillation sequence involving the 10-to-2 and/or 15-to-1 routines.

Table I shows the costs and improvements for a number of distillation sequences given an initial error probability of $p = 0.01$. Not surprisingly, the table shows that the 10-to-2 routine has a smaller cost, but the 15-to-1




FIG. 7. Log-log plot of the cost required to produce output states satisfying a goal error probability (e_g) given input states with error probability 0.01. The upper horizontal segments show the cost for the best sequence involving only the 15-to-1 routine. The lower horizontal segments show the cost for the best sequence of routines that achieves e_g or better. The gray regions indicate the improvement obtained using the 10-to-2 routine.

routine has greater improvement in error probability per round. For distillations that use both routines, we find numerically that if 15-to-1 rounds are used they should be placed first. This is intuitively consistent with the higher threshold for the 15-to-1 routine, which suggests better performance at high error probabilities.

The cost improvements shown in Tab. I are illustrated more visually in Fig. 7, which shows the production cost, at a fixed input error probability $p = 0.01$ and as a function of e_g , of the best distillation sequence compared to the best sequence using only the 15-to-1 routine. For example, a goal error probability near 10^{-5} can be achieved by using either two rounds of the 15-to-1 routine at a production cost of 261.7 or two rounds of the 10-to-2 routine at a cost of 27.9. In this case, the improvement in production cost obtained by incorporating the 10-to-2 routine is a factor of 9.4.

VI. CONCLUSIONS

Magic-state distillation enables universal quantum computing given only mediocre copies of a non-stabilizer state and high-quality Clifford operations. Considering the importance of Clifford-based techniques to the theory of fault tolerance, we expect that magic-state distillation will prove valuable for the practical implementation of quantum computers.

At the logical level, computationally useful quantum algorithms involve many non-Clifford gates, generally enough to account for a significant fraction of all gates

| Distillation scheme | Cost | Output error probability, $e(p)$ | Cost improvement factor |
|---------------------|--------|----------------------------------|-------------------------|
| <i>A</i> | 5.5 | 9×10^{-4} | 3.2 |
| <i>B</i> | 17.4 | 4×10^{-5} | 1 |
| <i>AA</i> | 27.9 | 7×10^{-6} | 9.4 |
| <i>BA</i> | 87.2 | 1×10^{-8} | 3.0 |
| <i>AAA</i> | 139.3 | 5×10^{-10} | 1.9 |
| <i>BB</i> | 261.7 | 2×10^{-12} | 1 |
| <i>BAA</i> | 436.2 | 1×10^{-15} | 9.0 |
| <i>AAAA</i> | 696.6 | 2×10^{-18} | 5.6 |
| <i>BBA</i> | 1308.7 | 2×10^{-23} | 3.0 |
| <i>BAAA</i> | 2180.8 | 1×10^{-29} | 1.8 |

TABLE I. Costs and output error probabilities at $p = 0.01$. The labels for the distillation schemes follow the convention given in Fig. 6. The cost improvement factor is with respect to the shortest sequence using only the 15-to-1 routine that achieves at least as good an output error probability.

employed. At least one high-quality magic state is required for the indirect implementation of each non-Clifford gate, so it is important to minimize the resources needed for the distillation of such states.

In this work, we contributed to the goal of resource reduction by introducing an $|H\rangle$ distillation routine that reduces the error probability for faulty $|H\rangle$ states from p to $O(p^2)$ and produces 2 output states using 10 input states. By judiciously combining this routine with the higher-order (p to $O(p^3)$) but higher-cost 15-to-1 routine

from Refs. [2, 3], we showed that the number of faulty $|H\rangle$ states required to distill states of a given quality can be reduced by up to an order of magnitude. Inclusion of additional distillation routines in the analysis would

likely lead to further improvements.

ACKNOWLEDGMENTS

We thank Scott Glancy for his help in bringing this work to fruition. This paper is a contribution of the National Institute of Standards and Technology and not subject to U.S. copyright.

-
- [1] Gabriele Nebe, E. M. Rains, and N. J. A. Sloane, “The invariants of the Clifford groups,” *Designs, Codes and Cryptography* **24**, 99–122 (2001), [arXiv:math/0001038](#).
 - [2] Sergey Bravyi and Alexei Kitaev, “Universal quantum computation with ideal Clifford gates and noisy ancillas,” *Phys. Rev. A* **71**, 022316 (2005), [arXiv:quant-ph/0403025](#).
 - [3] E. Knill, “Fault-tolerant postselected quantum computation: Schemes,” (2004), [arXiv:quant-ph/0402171](#).
 - [4] Ben Reichardt, “Quantum universality from magic states distillation applied to CSS codes,” *Quantum Information Processing* **4**, 251–264 (2005), [arXiv:quant-ph/0411036](#).
 - [5] Earl T. Campbell and Dan E. Browne, “On the structure of protocols for magic state distillation,” in *Lecture Notes in Computer Science: Theory of Quantum Computation* (2009) [arXiv:0908.0838](#).
 - [6] Earl T. Campbell and Dan E. Browne, “Bound states for magic state distillation in fault-tolerant quantum computation,” *Phys. Rev. Lett.* **104**, 030503 (2010), [arXiv:0908.0836](#).
 - [7] N. Cody Jones, Rodney Van Meter, Austin G. Fowler, Peter L. McMahon, Jungsang Kim, Thaddeus D. Ladd, and Yoshihisa Yamamoto, “A layered architecture for quantum computing using quantum dots,” (2010), [arXiv:1010.5022](#).
 - [8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2001).

Appendix

Fig. 8 provides additional details about the circuit identities used in Fig. 2.

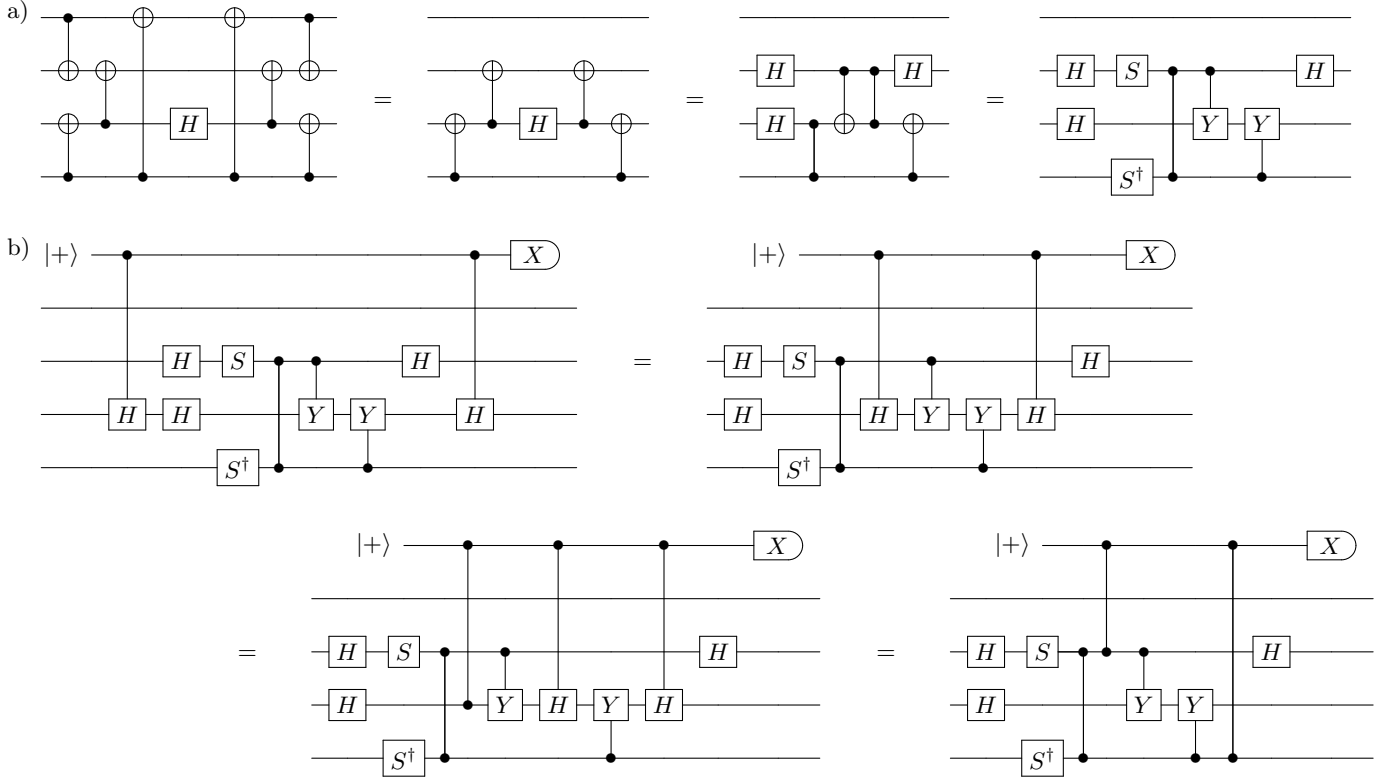


FIG. 8. Additional details regarding the relations between circuits in Fig. 2. (a) A sequence of circuit identities deriving the form of the encoded Hadamard gate used in Fig. 2(b). The starting circuit implements H on the second logical qubit of the code \mathcal{C}_4 by decoding the logical qubits into the first and third physical qubits, applying H to the third qubit, and re-encoding. The first equivalence is obtained by commuting and cancelling pairs of ${}^C X$ gates. The second equivalence uses the decomposition of the ${}^C X$ gate into ${}^C Z$ and H gates several times as well as the identity $H^2 = I$. The final equivalence uses two facts: $ZX = iY$ and the order of a ${}^C X$ and a ${}^C Z$ gate with the same target can be exchanged if a ${}^C Z$ gate is added between the controls. (b) A sequence of circuit identities showing why the pair of ${}^C H$ gates targeting the fourth qubit in Fig. 2(b) can be eliminated. Other than reorganization of commuting gates, these equivalences rely on the fact that, because H anticommutes with Y , ${}^C H$ and ${}^C Y$ gates with the same target can be exchanged if a ${}^C Z$ gate between the two controls is added.